

Detection and Characterization of Intrinsic Symmetry

ANIRBAN MUKHOPADHYAY and SUCHENDRA M. BHANDARKAR

Department of Computer Science, The University of Georgia, Athens, GA 30602-7404, USA

and

FATIH PORIKLI

Mitsubishi Electric Research Laboratories, Inc., 201 Broadway, Cambridge, MA 02139, USA

A comprehensive framework for detection and characterization of overlapping intrinsic symmetry over 3D shapes is proposed. To identify prominent symmetric regions which overlap in space and vary in form, the proposed framework is decoupled into a *Correspondence Space Voting* procedure followed by a *Transformation Space Mapping* procedure. In the correspondence space voting procedure, significant symmetries are first detected by identifying surface point pairs on the input shape that exhibit *local* similarity in terms of their intrinsic geometry while simultaneously maintaining an intrinsic distance structure at a *global* level. Since different point pairs can share a common point, the detected symmetric shape regions can potentially overlap. To this end, a *global intrinsic distance-based voting* technique is employed to ensure the inclusion of only those point pairs that exhibit significant symmetry. In the transformation space mapping procedure, the *Functional Map* framework is employed to generate the final map of symmetries between point pairs. The transformation space mapping procedure ensures the retrieval of the underlying dense correspondence map throughout the 3D shape that follows a particular symmetry. Additionally, the formulation of a novel cost matrix enables the inner product to successfully indicate the complexity of the underlying symmetry transformation. The proposed transformation space mapping procedure is shown to result in the formulation of a semi-metric symmetry space where each point in the space represents a specific symmetry transformation and the distance between points represents the complexity between the corresponding transformations. Experimental results show that the proposed framework can successfully process complex 3D shapes that possess rich symmetries.

Categories and Subject Descriptors:

General Terms: Symmetry, Intrinsic Geometry

Additional Key Words and Phrases: Intrinsic Symmetry, Biharmonic Distance, Symmetry Space, Symmetry Groups, Correspondence Space Voting, Transformation Space Mapping

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 0730-0301/YYYY/13-ARTXXX \$10.00

DOI XX.XXXX/XXXXXXXX.YYYYYYY

<http://doi.acm.org/XX.XXXX/XXXXXXXX.YYYYYYY>

1. INTRODUCTION

Symmetry is ubiquitous in nature and in artificial man-made objects. The detection and characterization of shape symmetry has attracted much attention in recent times, especially within the computer graphics community [MITRA13]. Although most of the existing literature has focused on the detection of *extrinsic* symmetries; a popular approach being transformation space voting [MITRA06]; there has been steadily growing interest in detection and characterization of *intrinsic* symmetries. Most recent efforts in intrinsic symmetry detection have focused on detection of global symmetries [OVSIJANIKOV08]; [LIPMAN10]. It is generally recognized that detection of overlapping intrinsic symmetry is a more challenging problem due to the larger search spaces involved in the detection of symmetric regions (in comparison to global symmetry analysis) and the determination of symmetry revealing transforms (in comparison to extrinsic symmetry detection). Also, overlapping intrinsic symmetry detection and characterization is more important because it is a more generalized problem in nature than extrinsic symmetry detection problem which can be considered as a special case of overlapping intrinsic symmetry detection.

We present a formal definition of overlapping intrinsic symmetry. An intrinsic symmetry over a shape is a subregion with associated self-homeomorphisms that preserve all pairwise intrinsic distances [MITRA13]. In this paper, we address the problem of intrinsic symmetry detection and characterization of shapes based on their intrinsic symmetries. Complex shapes often exhibit multiple symmetries that overlap spatially and vary in form as depicted in Figure 1. Overlapping symmetry analysis enables the construction of high-level representations that enhance the understanding of the underlying shape and facilitate solutions to such problems as shape correspondence, shape editing, and shape synthesis [WANG11]. However, analysis of overlapping symmetry poses additional challenges as described below.

Existing approaches to intrinsic symmetry detection, including those based on region growing [XU09], partial matching [RAVIV10], and symmetry correspondence [LIPMAN10] are not able to extract physically overlapping symmetries. Lipman et al. [LIPMAN10] cluster sample surface points from an input shape and use a symmetry correspondence matrix (SCM) to identify intrinsic symmetry properties of groups of surface points. In their approach, each SCM entry measures how symmetric two surface points are based on some measure of intrinsic geometric similarity between the local neighborhoods of the points. Xu et al. [XU12] let surface point pairs vote for their partial intrinsic symmetry and perform intrinsic symmetry grouping using a 2-step spectral clustering procedure. However, their approach lacks the ability to retrieve the final symmetry map which makes characterization of the specific intrinsic symmetry a difficult problem.

Our key idea behind our approach to intrinsic symmetry detection and characterization is to approach the problem from a shape




Fig. 1. Symmetry extraction in functional space.

correspondence perspective and generate the transformation map which can be further used to describe the symmetry space. To this end, we perform two stages of processing where in the first stage, representative symmetric point pairs are identified based on their local geometry and a global distance representation and in the second stage the original transformation is retrieved as a *map* to facilitate further characterization of the underlying symmetry. The detected intrinsic symmetries are quite general, as shown in Figure 1.

1.1 Symmetric point pairs

A fundamental question in symmetry detection is quantifying the extent of symmetry between a pair of points. The primary challenge in identifying potentially symmetric point pairs is to come up with conditions strong enough to adequately constrain the symmetry search space such that the symmetry detection procedure is computationally tractable. We rely on a local criterion, such as geometric similarity, and a global criterion, such as the distance-based symmetry support received by a point pair $\{a, b\}$, to detect and quantify the extent of symmetry between points a and b . We refer to a point pair $\{a, b\}$ which satisfies the local geometric similarity criterion as a *good voter*. The point pairs within the population of *good voters*, that enjoy sufficiently strong global distance-based symmetry support are deemed to be *symmetric* point pairs. The global symmetry support for point pair $\{a, b\}$ is quantified by the number of other point pairs which potentially share the intrinsic symmetry properties of $\{a, b\}$. The global symmetry support is computed using a simple, distance-based symmetry criterion defined over two point pairs within the population of *good voters* as shown in Figure 2.

1.2 Overview

The input to the proposed symmetry detection and characterization algorithm is a 3D shape that is approximated by a 2-manifold triangular mesh. The local intrinsic geometry is quantified using the *wave kernel signature* (WKS) [AUBRY11] whereas the global intrinsic geometric distance measure chosen is the *biharmonic distance measure* (BDM) [LIPMAN10a]. The proposed algorithm consists of a voting procedure on the correspondence space defined over a set of locally symmetric surface point pairs sampled from the input 3D shape, followed by the generation of a functional map (see Figure 2).

1.2.0.1 Correspondence Space Voting. The input to the correspondence space voting procedure comprises of point pairs that can be considered as candidates for symmetry. Prior to performing the transformation space mapping procedure, we sample a set of locally symmetric point pairs from the input shape based on the similarity of their WKS's. To estimate the symmetry support received by a point pair, we perform a voting procedure by counting the number of *good voters* which potentially share the same intrinsic symmetry as the source point pair. The voting procedure ensures that we have a set of good point pair initializations from which we can create an initial map of the symmetry transformation that can then be extrapolated to other surface point pairs on the 3D shape.

1.2.0.2 Computation of the Functional Map. A functional map provides an elegant representation for the maps between surfaces, allowing for efficient inference and manipulation. In the functional map approach, the concept of a map is generalized to incorporate correspondences between real valued functions rather than simply between surface points on the 3D shapes. Our choice of the multi-scale eigenbasis of the Laplace-Beltrami operator makes

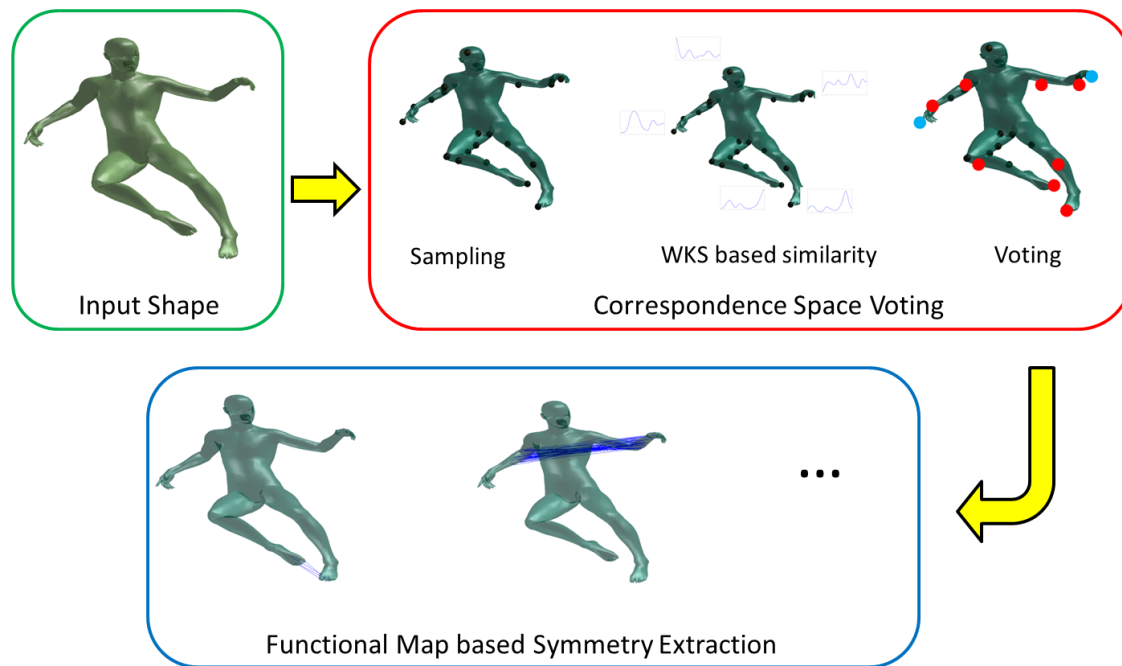


Fig. 2. Overview of the proposed symmetry detection and characterization framework.

the functional map representation both, very compact and yet informative. We show that our functional map formulation not only results in a compact description of the underlying map of the symmetry transformation, but also enables meaningful characterization of the symmetry transformation.

The primary contributions of our paper can be summarized as follows:

- (1) Robust and meaningful characterization of the symmetry transformation via formulation of a symmetry space to quantitatively distinguish between instances of simple and complex symmetry.
- (2) Leveraging of the functional map representation to successfully represent the map of the underlying symmetry transformation regardless of its complexity.
- (3) Providing a solution to the symmetric flipping issue via interpolation of the correspondence from the base surface point pairs using functional maps.
- (4) Enabling recovery of the symmetry groups via clustering on the functional maps.

2. RELATED WORK

The research literature on symmetry detection has grown substantially in recent years as shown in the excellent survey by Mitra et al. [MITRA13]. In this paper, we do not attempt to provide an *exhaustive* exposition of the state of the art in symmetry detection; rather we focus on discussing existing works that are most closely related to our proposed approach.

2.1 Symmetry detection in transformation space

Several recent approaches to detect approximate and partial extrinsic symmetries have focused on algorithms that cluster

votes for symmetries in a parameterized transformation space [IMIYA99]; [MITRA06]; [YIP00]; [LI05]]. Mitra et al. [MITRA06] generate votes in a transformation space to align pairs of similar points and then cluster them in a space spanned by the parameters of the potential symmetry transformations. Regardless of how good the shape descriptors are, the aforementioned methods are not effective at finding correspondences between points in complex symmetry orbits that are spread across multiple distinct clusters in the transformation space. Since the dimensionality of the transformation space increases with the complexity of the symmetry, the voting procedure in transformation space becomes increasingly intractable when dealing with complex symmetries.

2.2 Symmetry representation in transformation space

There exists a body of published research literature that characterizes shape representations based on the extent of symmetry displayed by an object with respect to multiple transformations. Kazdhan et al. [KAZDHAN03] have proposed an extension of Zabrodskys symmetry distance to characteristic functions, resulting in a *symmetry descriptor* that measures the symmetries of an object with respect to all planes and rotations through its center of mass. Podolak et al. [PODOLAK06] have extended the symmetry descriptor to define a planar reflective symmetry transform (PRST) that measures reflectional symmetries with respect to all planes through space. Rustamov et al. [RUSTAMOV08] have extended the PRST to consider surface point-pair correlations at multiple radii. Although the above representations provide a measure of symmetry for a regularly sampled set of transformations within a group, they are practical only for transformation groups of low dimensionality (for example, rigid body transformations would require one to store

a six-dimensional matrix) and break down when faced with groups of higher dimensionality.

2.3 Discovery of repeating structures

There exists a class of techniques that exploits the redundancy in repeating structures to robustly detect symmetries [BOKELOH09]; [LI06]; [LIU07]; [PAULY08]; [SHIKHARE01]. The transformation space voting method proposed by Mitra et al. [MITRA06] is extended in [PAULY08] by fitting parameters of a transformation generator to optimally register the clusters in transformation space. Berner et al. [BERNER08] and Bokeloh et al. [BOKELOH09] have taken a similar approach using subgraph matching of feature points and feature lines, respectively, to establish potential correspondences between repeated structures. This is followed by an iterative closest points (ICP) algorithm to simultaneously grow corresponding regions and refine matches over all detected patterns, allowing the detection of repeated patterns even in noisy data [PAULY08], but at the cost of requiring a-priori knowledge of the commutative group expected in the data. Also, the non-linear local optimization procedure within the ICP algorithm could cause it to get trapped in a local minimum if the initialization is not good enough.

2.4 Eigen-analysis methods

Lipman et al. [LIPMAN10] have proposed an eigen-analysis technique for symmetry detection that relies on spectral clustering. The top eigenvectors of their geometric similarity-based SCM characterize the symmetry-defining orbits, where each orbit includes all points that are symmetric with one another. However, their work is not suited for multi-scale partial symmetry detection. First, expressing local point similarities as symmetry invariants is only appropriate for global intrinsic symmetry detection. In the case of partial symmetry detection, it is not always possible to reliably judge if two surface points are symmetric by comparing only their point (i.e., local) signatures, especially when one point lies on the boundary of symmetric regions. Moreover, their single-stage clustering procedure is unable to identify overlapping symmetries. Xu et al. [XU12] have extended the eigen-analysis approach of Lipman et al. [LIPMAN10] by incorporating the concept of global intrinsic distance-based symmetry support accompanied by a 2-stage spectral clustering procedure to distinguish between scale detection and symmetry detection. Although they showed some interesting results, the 2-stage spectral clustering procedure made their method extremely slow. Furthermore, the absence of transformation map retrieval meant that further processing of the detected symmetries, which are represented as point pairs, was extremely inefficient.

The proposed scheme is decoupled into two steps of *Correspondence Space Voting* and *Transformation Space Mapping*. The *Correspondence Space Voting* technique is inspired from the work of Xu et al. [XU12], but in our technique, we bypassed the 2 particularly lengthy steps of spectral clustering and the all-pair geodesic distance calculation to improve the running time quite significantly. Moreover, our introduction of *Transformation Space Mapping* in symmetry detection is quite novel in the sense that this not only provides a concise description of the underlying symmetry transformation, but according to our knowledge this is one of the first works which has the unique ability of characterizing the symmetric transformation.

3. THEORETICAL FRAMEWORK

In this section, we present a formal description of the problem being addressed in this paper. In particular, we define the input to and output of the proposed algorithm. We also define the type of intrinsic symmetries the proposed algorithm is designed to detect and the formal characterization of these symmetries. There are two primary aspects to the theoretical framework for the proposed algorithm, i.e., *correspondence space voting* (CSV) and *functional map retrieval* (FMR). In the case of CSV, a joint criterion, that combines local intrinsic surface geometry and global intrinsic distance-based symmetry, is proposed and shown to result in a provably necessary condition for intrinsic symmetry. In the case of FMR, a formal scheme for the characterization of the detected symmetries, based on their complexity is proposed.

We have mainly restricted our study of intrinsic symmetries to isometric involutions for two reasons; first, having a provably necessary condition for intrinsic symmetry provides theoretical soundness and second, the proposed symmetry criterion bounds the search space sufficiently, ensuring that the solution is computationally tractable [XU12]. Following the initial detection of good symmetric correspondences, we proceed to retrieve the map of the symmetry transformation by leveraging the functional map framework. Moreover, the intrinsic symmetry ensures that the retrieved maps exhibit a diagonality characteristic. A cost matrix is designed to exploit this characteristic such that the inner product of functional map with the cost matrix results in a *quantitative* evaluation of the complexity of the detected symmetry.

3.1 Input to the Proposed Algorithm

Our problem domain is a compact, connected 2-manifold, M , with or without a boundary. The manifold M is a *3D shape*, i.e., $M \rightarrow \mathbb{R}^3$. Distances on the manifold M are expressed in terms of an intrinsic distance measure. In particular, the biharmonic distance measure (BDM) is used on account of its ease of calculation and greater robustness to local surface perturbation when compared to the geodesic distance measure d_M [LIPMAN10a]. In the remainder of the paper, we use the term intrinsic distance and biharmonic distance interchangeably.

3.2 Output of the Proposed Algorithm

In the proposed algorithm each symmetry transformation is represented by a map. Consequently, the output of the proposed algorithm consists of maps represented as matrices. The output can be regarded as a complete description of all the overlapping intrinsic symmetries represented in a compact and informative manner.

3.3 Definition of Intrinsic Symmetry

Suppose we are given a compact manifold M without boundary. Following [RAVIV07] we call M intrinsically symmetric if there exists a homeomorphism $T: M \rightarrow M$ on the manifold that preserves all geodesic distances. That is:

$$d_M(p, q) = d_M(T(p), T(q)) \forall p, q \in M \quad (1)$$

where $d_M(p, q)$ is the intrinsic distance between two points on the manifold. In this case, we call the mapping T an intrinsic symmetry.




Fig. 3. Limitation of the proposed symmetry detection algorithm.

3.4 Symmetry Criteria

We propose two simple criteria to test whether two surface point pairs on the manifold potentially share the same intrinsic symmetry. Specifically, given two surface point pairs $\{x, x'\}$ and $\{y, y'\}$ on manifold M , the first criterion, which is based on local intrinsic geometry, determines the symmetry potential of the two surface point pairs by comparing their corresponding wave kernel signatures (WKS's) as follows:

$$\begin{aligned} WKS(x, t) &= WKS(x', t) \text{ and} \\ WKS(y, t) &= WKS(y', t) \forall t \geq 0 \end{aligned} \quad (2)$$

where t is a scale parameter. The second criterion is based on intrinsic distance as follows:

$$\begin{aligned} d_M(x, y) &= d_M(x', y') \text{ and} \\ d_M(x, y') &= d_M(y, x') \end{aligned} \quad (3)$$

The above two criteria are necessarily satisfied if the surface point pairs under consideration correspond to the same intrinsic symmetry.

3.5 Limitations

The proposed framework fails to detect intrinsic symmetry in cases where the second symmetry criterion (in equation (3)) is not satisfied. For example, Figure 3, depicts two human figures that form what may be perceived as a translational symmetry, although they do not possess intrinsic symmetry. The second symmetry criterion (equation (3)) fails to hold in this situation. Generally speaking, the proposed algorithm is not designed to detect all forms symmetry resulting from repeated patterns, especially if the patterns are not connected. In contrast to the approach of Xu et al. [XU12], the use of biharmonic distance instead of geodesic distance as the intrinsic distance measure ensures that the proposed algorithm is capable of detecting intrinsic symmetry even in the presence of small perturbations (such as small bumpy regions) on the 3D surface [LIPMAN10a].

4. CORRESPONDENCE SPACE VOTING

The first step in the proposed symmetry detection and characterization framework is the correspondence space voting (CSV) procedure. Although a voting procedure has been previously incorporated in an earlier symmetry detection technique, it was carried out primarily in transformation space [MITRA06]. The importance of correspondence space for the detection of symmetry was explained more recently by Lipman et al. [LIPMAN10]. In our case, although



Fig. 4. Sampling results on three different human shapes

the detected symmetry is finally represented in functional space, it is critical to have good initialization to ensure the success of the final map generation. We have designed and implemented a CSV algorithm to facilitate good initial guesses. The CSV algorithm comprises of three stages described in the following subsections:

4.1 Point Selection

A subset of points with adequate discriminative power needs to be selected prior to the generation of surface point pairs. A subset X consisting of n sample points is chosen from the surface of the given input 3D shape using the *Farthest Point Sampling* strategy [ELDAR97]. Although originally designed to operate on geodesic distances generated using the marching cubes algorithm, in our case we have employed the farthest point random sampling strategy in biharmonic distance space. The results of the sampling procedure are depicted in Figure 4.

The *biharmonic distance measure* (BDM) is similar in form to the *diffusion distance measure* (DDM) [LING06]. The BDM kernel is based on the Green's function of the biharmonic differential equation. In the continuous case, the (squared) biharmonic distance between two points x and y can be defined using the eigenvectors (ϕ_k) and eigenvalues (λ_k) of the Laplace- Beltrami operator [LIPMAN10] as follows:

$$d_M(x, y)^2 = \sum_{k=1}^{\infty} \frac{(\phi_k(x) - \phi_k(y))^2}{\lambda_k^2} \quad (4)$$

The above definition of the BDM is slightly different from that of the DDM where the denominator in the case of the DDM is $e^{2t\lambda_k}$. However, this subtle change ensures greater control over the characterization of the global and local properties of the underlying manifold in the case of the BDM. Consequently, the BDM is fundamentally different from the DDM with significantly different properties.

The BDM, as expressed in equation (4), captures the rate of decay of the normalized eigenvalues λ_k of the Laplace- Beltrami operator; if the decay is too slow, it produces a logarithmic singularity along the diagonal of the Green's function [YEN07]. Alternatively, too fast a decay basically ignores eigenvectors associated with higher frequencies, resulting in the BDM being global in nature (i.e., the local surface details are ignored). Lipman *et al.* [LIPMAN10a], demonstrated that performing quadratic normalization provides a good balance, ensuring that the decay is slow enough to capture the local surface properties around the source point and yet rapid enough to encapsulate global shape information.

In particular, Lipman *et al.* [LIPMAN10a] have theoretically proved two important properties of the BDM, i.e., that it is (i) a metric, and (ii) smooth everywhere except at the source point where it is continuous. The key observation is that for 3D surfaces, the eigenvalues λ_k , $k = 1, 2, \dots$, of the Laplacian are an increasing function of k resulting in the continuity of the BDM everywhere

and also smoothness of the BDM everywhere except at the source point, where it has only a derivative discontinuity.

In our implementation of the farthest point random sampling strategy, a single point is selected randomly at first and the remaining points are chosen iteratively from remainder set by selecting the farthest point in the biharmonic distance space at each iteration. This strategy generates a set of points located mostly in the vicinity of the shape extrema which can then be used in the subsequent surface point pair generation procedure.

4.2 Surface Point Pair Generation

From the chosen subset X consisting of n sample points, the surface point pairs are generated by exploiting the similarity of their local intrinsic geometric structure. The similarity of local intrinsic geometric structure of two surface points is determined by comparing their corresponding wave kernel signatures (WKS's). To determine the WKS of a surface point x , one evaluates the probability of a quantum particle with a certain energy distribution to be located at point x . The behavior of the quantum particle on the surface is governed by the Schrodinger equation. Assuming that the quantum particle has an initial energy distributed around some nominal energy with a probability density function $f(e)$, the solution of the Schrodinger equation can then be expressed in the spectral domain as:

$$\psi(x, t) = \sum_{k>=1} e^{ie_k t} f(e_k) \phi_k(x) \quad (5)$$

Aubry et al. [AUBRY11] considered a family of log-normal energy distributions centered around some mean log energy $\log e$ with variance σ^2 . This particular choice of distributions is motivated by a perturbation analysis of the Laplacian spectrum [AUBRY11]. Having fixed the family of energy distributions, each point x on the surface is associated with a WKS of the form: $p(x) = (p_{e_1}(x), \dots, p_{e_n}(x))^T$ where $p_{e_i}(x)$ is the probability of measuring a quantum particle with the initial energy distribution $e_i(x)$ at point x . Aubry et al. use logarithmic sampling to generate the values $e_1(x), \dots, e_n(x)$.

The WKS can be shown to exhibit a band-pass characteristic. This reduces the influence of low frequencies and allows better separation of frequency bands across the descriptor dimensions. As the result, the wave kernel descriptor exhibits superior feature localization compared to the Heat Kernel Signature (HKS) [LITMAN13]. Although the WKS is invariant under isometric deformation, in most practical cases where the underlying surface is represented by a discrete triangular mesh, it is not possible to strictly satisfy the invariance criterion in equation (2). Consequently, we have considered the 2-norm (simply squared Euclidian distance in the WKS space) of the WKS and chosen surface point pairs from the subset X which satisfy the WKS invariance criterion in equation (2) to within a prespecified threshold T_{WKS} instead of requiring strict equality. The surface point pairs that satisfy the WKS invariance criterion in equation (2) are considered for the next step of global distance-based voting. This relaxation of the WKS invariance criterion ensures that even the overlapping intrinsic symmetries are considered for voting. We have coined the term *good voters* to denote the subset of surface point pairs that satisfy the WKS invariance criterion (to within the prespecified threshold).

4.3 Global Distance-based Voting

The global distance-based voting step in the proposed symmetry detection technique is inspired by the work of Xu et. al. [XU12]. A subset of symmetric point pairs is extracted from the set of *good*

voters using the global distance-based voting procedure prior to functional map generation. The goal of the voting procedure is to accumulate symmetry support for the *good voters* and extract the symmetric point pairs based on the level of symmetry support received.

A point pair in the set of *good voters* is deemed to be symmetric if it has a sufficiently large global symmetry support, which in [XU12] is measured by the number of point pairs that satisfy the intrinsic distance criterion in equation (3). Xu et al. [XU12] have presented a voting technique based on two point pairs followed by spectral clustering on the symmetric point pairs which enables one to distinguish whether a point pair supports one particular symmetry or more than one type of symmetry. However, since we are not interested, at this stage, in making the above distinction between the point pairs, we have chosen to adopt a straightforward approach of choosing a point pair and letting the set of *good voters* vote and decide whether or not the particular point pair satisfies the intrinsic distance criterion in equation (3). Another modification to the previous voting procedure of Xu et al. [XU12] is our use of the biharmonic distance as the intrinsic distance instead of the geodesic distance. This modification provides two advantages over the previous voting procedure. First, since the computation of geodesic distances between all surface point pairs is more intensive than the computation of biharmonic distances, our procedure is much faster. Second, biharmonic distances are less sensitive to noise and surface perturbations compared to geodesic distances, making our procedure more robust.

5. SYMMETRY EXTRACTION USING FUNCTIONAL MAPS

Before describing our functional map approach to symmetry extraction, we first provide an overview of the functional map framework proposed by Ovsjanikov et al. [OVSJANIKOV12]. A functional map is a novel approach for inference and manipulation of maps between shapes that tries to resolve the issues of correspondences in a fundamentally different manner. Rather than plotting the corresponding points on the shapes, the mappings between functions defined on the shapes are considered. This notion of correspondence generalizes the standard point-to-point map since every pointwise correspondence induces a mapping between function spaces, while the opposite, in general, is not true.

The new framework described above provides an elegant way to avoid direct representation of correspondences as mappings between shapes using a functional representation. Ovsjanikov et al. [OVSJANIKOV12] have noted that when two shapes X and Y are related by a bijective correspondence $t : X \rightarrow Y$, then for any real function $f : X \rightarrow \mathbb{R}$, one can construct a corresponding function $g : Y \rightarrow \mathbb{R}$ as $g = f \circ t^{-1}$. In other words, the correspondence t uniquely defines a mapping between the two function spaces $F(X, \mathbb{R}) \rightarrow F(Y, \mathbb{R})$, where $F(X, \mathbb{R})$ denotes the space of real functions on X . Equipping X and Y with harmonic bases, $\{\phi_i\}_{i \geq 1}$ and $\{\psi_j\}_{j \geq 1}$, respectively, one can represent a function $f : X \rightarrow \mathbb{R}$ using the set of (generalized) Fourier coefficients $\{a_i\}_{i \geq 1}$ as $f = \sum_{i \geq 1} a_i \phi_i$. Translating this representation into the other harmonic basis $\{\psi_j\}_{j \geq 1}$, one obtains a simple representation of the correspondence between the shapes given by $T(f) = \sum_{i, j \geq 1} a_i c_{ij} \psi_j$ where c_{ij} are Fourier coefficients of the basis functions of X expressed in the basis of Y , defined as $T(\phi_i) = \sum_{j \geq 1} c_{ij} \psi_j$. The correspondence t between the shapes can thus be approximated using k basis functions and encoded using a $k \times k$ matrix $C = (c_{ij})$ of these Fourier coefficients, referred to as the functional matrix. In this representation, the computation




Fig. 5. The band-pass characteristic of the WKS at different scales for the same *human* figure.

of the shape correspondence $t : X \rightarrow Y$ is translated into a simpler task of determining the functional matrix C from a set of correspondence constraints. The matrix C has a diagonal structure if the harmonic bases $\{\phi_i\}_{i \geq 1}$ and $\{\psi_j\}_{j \geq 1}$ are compatible, which is a crucial property for the efficient computation of the correspondence.

In our symmetry extraction algorithm, instead of comparing two different shapes, we propose to compare two symmetric regions within the same shape. In particular, based on the previously detected set of symmetric point pairs, we leverage the functional map representation in the following manner: for each pair of symmetric points, we deem one point as the source and the other as the destination and choose a local region around each point. The ordering of the source and destination points within the pair is the same as originally chosen during the voting procedure. The corresponding eigenbases for the points in the source and destination regions are computed. These eigenbases are ordered based on their similarity with each other and the final functional map for that particular symmetry is extracted. The functional map representation ensures that (a) the problem of symmetry extraction is tractable and, (b) the resulting symmetry can be represented, not by a large matrix of point correspondences, but rather as a more compact map which can be further manipulated for other applications as well.

6. SYMMETRY CHARACTERIZATION USING FUNCTIONAL MAPS

In general, the characterization of a specific transformation based on its functional map is a challenging task. However in our case, the proposed CSV framework ensures that the point pairs used in the generation of the functional map are intrinsically symmetric to a reasonable extent. This property of intrinsic symmetry ensures that the resulting functional map is diagonal or close to diagonal [OVSEJANIKOV12]. However, in reality, there are several cases where the actual transformation deviates substantially from ϵ -isometric deformation resulting in a densely populated functional matrix C . The diagonality property of the C matrix was first exploited successfully in an intrinsic correspondence framework using a sparse modeling technique [PORKASS13]. In their paper, the authors introduced a weight matrix W of same size as C where lower weights are assigned to elements close to the diagonal, and larger weights are assigned to the farther off-diagonal elements by using an inverted Gaussian model where the 0 weights are assigned at diagonal and the weights at off-diagonals are larger. The element-wise multiplication of W with C is a determining factor in the assessment of the diagonality of C since the matrix inner product results in higher values for the off-diagonal elements and lower values for the diagonal elements thereby resulting in a measure of the diagonality of the matrix C .




Fig. 6. Symmetry characterization based on functional maps, in particular, the complexity of the symmetry transformation is characterized by the weight matrix W and represented in the increasing order of the value of inner product of W and C .

In the context of symmetry characterization problem, we assume that the off-diagonality of C corresponds to the complexity of the symmetry transformation and therefore is penalized during multiplication with elements of the weight matrix W , i.e., more non-zero off-diagonal elements in C would represent a more complex symmetry transformation. Thus, we can, not only determine the complexity of the symmetry transformation but can also successfully formulate a 1D semi-metric symmetry space, wherein each symmetry transformation is represented as a point in the symmetry space with a value given by the inner product between the matrix C and weight matrix W . The Euclidian distance between the points in the 1D symmetry space represents the complexity distance between the transformations. In the symmetry space, any perfectly isometric transformation will result a point at origin of the line. On the other hand, more complex symmetric transformations would result in points farther away signifying more complex transformations. This space is called a semi-metric space because of it follows two main properties of the distance definition but not the third one. For two points X and Y in the symmetry space, the distance $d(X, Y)$ from X to Y always follows

- (1) Non-negativity: $d(X, Y) \geq 0$
- (2) Symmetry: $d(X, Y) = d(Y, X)$
but not
- (3) Triangle Inequality.

It is also possible to cluster the points in the 1D symmetry space to identify intrinsic symmetries which are potentially similar in nature as shown in the Experimental Results section to follow.

7. EXPERIMENTAL RESULTS

In this section, we present and discuss the results obtained by the proposed intrinsic symmetry detection algorithm on 3D shapes. We also provide comparisons of our results with those obtained from the most closely related approaches [XU12], [LIPMAN10]. We also shown some applications where the detected symmetries can be further analyzed for symmetry characterization and clustering, potentially revealing greater semantic information about the underlying 3D shape. Most of the 3D shape models used in our experiments are from the *Non-rigid World* dataset [BRONSTEIN07] unless mentioned otherwise.

7.1 Preprocessing

Several discrete schemes have been proposed in recent years to approximate the Laplace-Beltrami operator on triangular meshes. Among these, the one mostly used most widely for computing the discrete Laplace operator is the cotangent (COT) scheme, originally proposed by Pinkall and Polthier [PINKALL93]. Recently, Belkin et al [BELKIN08] proposed a discrete scheme based on the heat equation, which has been proven to possess the point-wise convergence property for arbitrary meshes. Although it is well known that no discrete Laplace operator can share all of the properties of its continuous counterpart, in our experiments the aforementioned discrete schemes produce eigenfunctions that approximately preserve the convergence property of the continuous Laplace operator for a reasonably well sampled triangular mesh. In all of the experiments described below, we have used the cotangent scheme [PINKALL93] for the computation of the discrete Laplace-Beltrami operator.

7.2 Symmetry Detection

The results of the proposed symmetry detection algorithm are depicted in Figures 1, 7, 8 and 12. Several important properties of the proposed algorithm are highlighted in these results.

7.2.1 General Symmetry Detection. The ability of the proposed algorithm to identify multiple intrinsic symmetries is evident from the results shown in Figures 7 and 12. The extracted symmetries are seen to cover the global symmetry of the underlying 3D shape which has undergone approximate isometric deformations. Additionally, the proposed algorithm is also observed to be capable of detecting symmetry transformations that cover individual components of a 3D object that possess various forms of self-symmetry.

7.2.2 Overlapping Symmetry Detection. One particularly important aspect of the proposed algorithm is its ability to detect instances of overlapping symmetry. An instance of overlapping symmetry is deemed to occur when a specific region on the surface of the 3D shape is simultaneously subjected to more than one symmetry transformation and, as a result, is symmetric to more than one region on the 3D shape surface. For example, Figure 7 shows that the overlapping symmetry between *all* paws of the *Cat* shape model is successfully detected by the proposed technique. Six different combinations of symmetry transformation between the four paws are depicted in Figure 7.

7.2.3 Noise. Although the CSV procedure is statistical in nature, the proposed biharmonic distance-based voting scheme ensures its robustness to noise. In particular, Figure 8 demonstrates the robustness of the proposed symmetry detection technique for different levels of synthetic Gaussian noise added to the *Human* shape model.




Fig. 7. Overlapping symmetry detection on the *Cat* shape model.

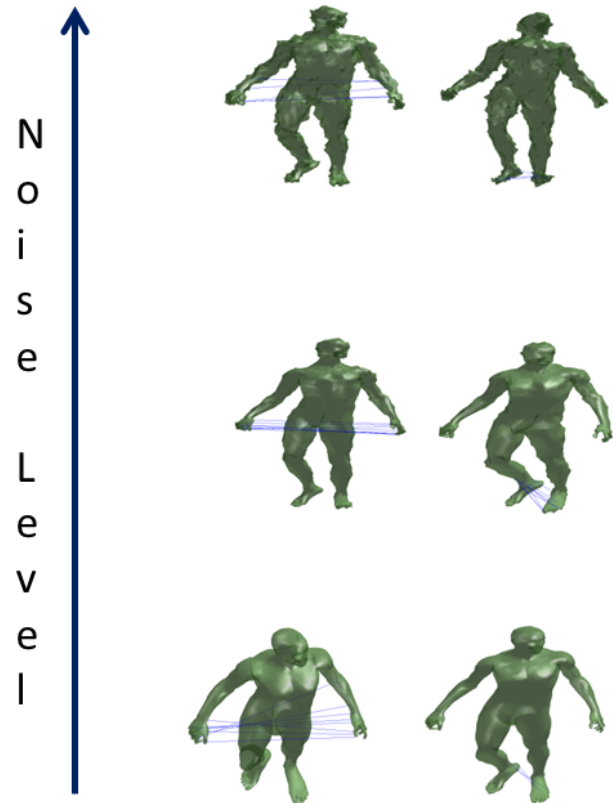


Fig. 8. Robustness of symmetry detection for three different levels of synthetic white Gaussian noise added to the *Human* shape model.

Table I. Timing results for the various steps in the proposed algorithm.

Points	Biharmonic	FMap	Extraction	Total
3400	32	2	5	39
10000	41	3	7	51
50000	237	4	15	256

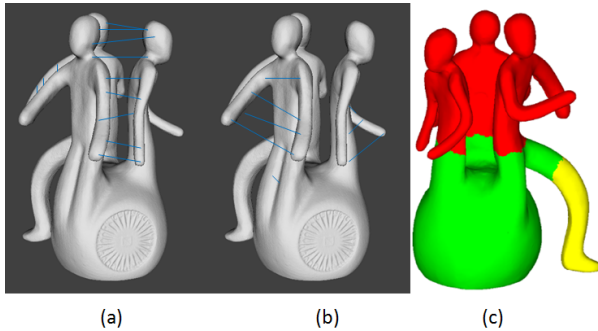


Fig. 9. In comparison of (c) the symmetry-factored embedding (SFE) technique of Lipman et al. [LIPMAN10], which primarily detects instances of global intrinsic non-overlapping symmetry, the proposed CSV procedure ensures the detection of instances of overlapping intrinsic symmetry (a,b) as well.

7.3 Performance Statistics

All of the experiments reported in this paper were performed on an Intel Core™ 3.4 GHz machine with 8 GB RAM. For all the example models, the number of sample points were in the range [20, 100]. Table I reports the timing results for the various steps in the proposed symmetry detection algorithm. In particular, unlike [XU12], wherein the most time consuming step of all-pairs geodesic distance computation is not reported, we also report the timing results for the equivalent step in our formulation i.e., the all-pairs biharmonic distance computation. The most time consuming step in the proposed algorithm, i.e., the all-pairs biharmonic distance computation, accounts for around 80% of the execution time of the proposed algorithm. More importantly, bypassing the two-step spectral clustering procedure described in [XU12] reduces significantly the computation time of the proposed algorithm.

7.4 Comparisons

We have compared the proposed symmetry detection algorithm primarily with methods that could be deemed sufficiently similar [XU12], [LIPMAN10]. The symmetry detection technique of Xu et al. [XU12] can detect overlapping partial intrinsic symmetries whereas that of Lipman et al. [LIPMAN10] is designed to deal with partial extrinsic symmetries. However, if the symmetric sub-shapes do not undergo significant pose variations, the global alignment component of [LIPMAN10] may allow it to detect certain partial intrinsic symmetries as well. However, whereas both methods [XU12], [LIPMAN10] are capable of detecting instances of partial intrinsic symmetry, neither is able to characterize the underlying symmetry.

In contrast, the proposed algorithm, not only detects overlapping intrinsic symmetries, but it also has the ability to characterize and cluster the detected symmetries in symmetry space. Since their entire formulation is based solely on global intrinsic distance-based voting, the technique of Xu et al. [XU12] suffers from the short-

coming of not being able to reliably detect symmetry flips. The symmetry flip phenomenon is deemed to occur when for the same symmetric transform, the point correspondences interchange their relative positions due to ill-posed symmetry detection criteria. The proposed algorithm, on the other hand, interpolates the functional map of symmetry transformations from the chosen directed point pair, to the remaining point pairs, i.e., once the source and destination within a point pair are identified, the remaining correspondences are obtained via interpolation using the functional map. The functional map ensures that the interpolated point pairs observe the same direction of symmetry as the initial directed pair, i.e., the functional map preserves the relative positions of points within a point pair that are subject to the same symmetry transformation. The functional map thus provides a natural solution to the symmetry flip problem as evidenced in Figure 10. The symmetry-factored embedding (SFE) technique of Lipman et al. [LIPMAN10], though designed for extrinsic symmetry detection, in simpler cases, is able to detect intrinsic symmetries as well. However, it fails completely in cases of overlapping symmetries whereas the proposed CSV procedure ensures the detection of instances of overlapping intrinsic symmetry as shown in Figure 9.

7.5 Quantitative Evaluation

In this section, we present quantitative evaluation of the results obtained by the proposed approach. In order to evaluate the proposed algorithm, we used the SHREC 2010 feature detection and description benchmark [BRONSTEIN10]. This dataset comprises of three shapes (null shapes) and a set of shapes obtained by applying a set of transformations on the null shapes. Shapes have approximately 10,000 to 50,000 vertices and their surfaces are represented by triangular meshes. We have specifically considered shapes which have undergone changes characterized by isometry, topology, micro-holes, scale and noise. For each transformation, in the initial phase, a total 50 sample points are generated using the farthest random sampling strategy as shown in Figure 11. Almost all transformations maintain a high repeatability (greater than 80%) at overlap values ≤ 0.75 except for changes in topology.

7.6 Symmetry Characterization

The diagonality property of the functional map could be used to characterize the underlying symmetry transformation and classify it as a simple transformation or as one that is more complex in nature. Our assumption is that, greater the complexity of symmetry transformation, greater the deviation of the shape deformation from intrinsic isometry, resulting in a deformation characterized by a functional matrix C with higher off-diagonal element values. The resulting characterization of the isometric deformation is depicted in Figure 6.

7.7 Symmetry Group Retrieval

As discussed in [XU12], symmetry has a group structure. Consequently, the retrieval of symmetry transformation can be cast as a clustering problem from an algorithmic perspective. We exploit the functional maps generated previously to cluster the detected instances of symmetry in transformation space and retrieve the symmetry groups. Due to the symmetric flip problem inherent in the CSV procedure, we adjusted the relative positions of source and destination points within each point pair by finding the minimum distance between the source and destination points in the biharmonic distance space. This ensures that the functional maps generated from potentially similar transformations will have similar

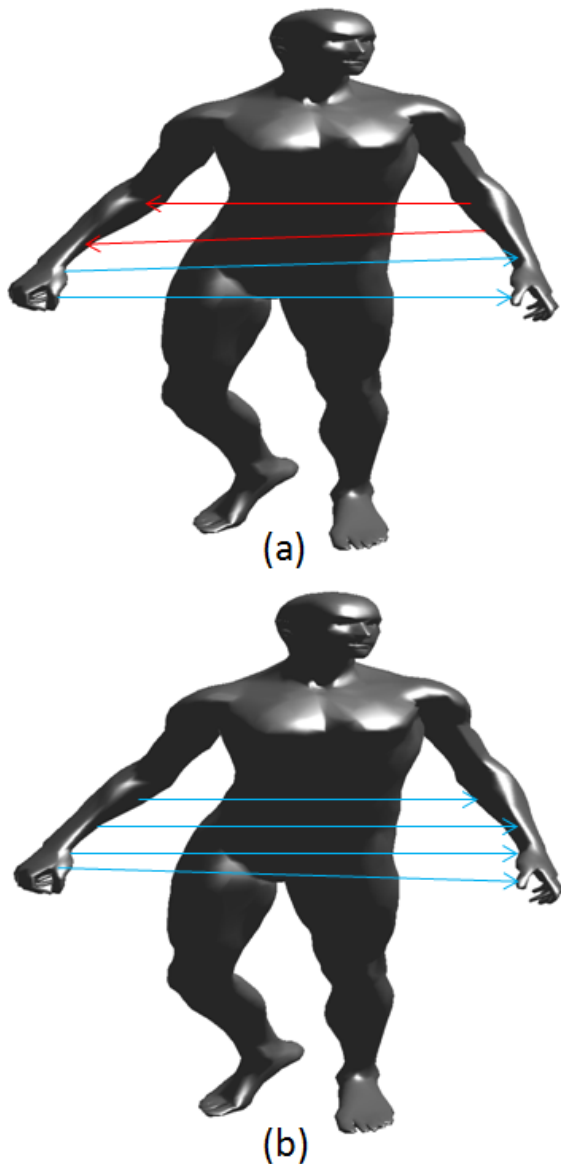


Fig. 10. In contrast to the technique of Xu et al. [XU12] (a), which relies solely on global intrinsic distance and consequently, fails to detect symmetric flips (in red), the proposed functional map-based interpolation of the symmetric correspondence (b) ensures robust detection of symmetric flips.

structure since the symmetry flip problem between the maps is resolved. In particular, we have used a simple k -means clustering algorithm to cluster the functional maps based on their symmetry groups as depicted in Figure 12.

8. CONCLUSION AND FUTURE DIRECTIONS

We have presented an algorithm for detection and characterization of intrinsic symmetry in 3D shapes. While the results obtained are encouraging, we regard our work as an initial attempt towards complete understanding of and a verifiable solution to the general problem of symmetry detection and characterization. We have identified




Fig. 11. Overlap vs. Repeatability plot for the proposed symmetry detection technique on the SHREC 2010 benchmark dataset.

some limitations of our approach and we hope to address these in our future work.

To the best of our knowledge, this is one of the first attempts to formalize the symmetry analysis problem not only as one of symmetry detection, but as one that can be extended to include symmetry characterization and symmetry clustering in the transformation space. In particular, the introduction of the functional map formalism in symmetry detection enables us to come up with a novel representation of the symmetry transformation as a map. In future, we aim to formulate operations, such as addition and subtraction, on these generated maps that would potentially provide a deeper and more comprehensive understanding of intrinsic symmetry in general.

The incorporation of *Transformation Space Mapping* in Symmetry characterization is a completely new idea and the full potential of it can only be realized after more extensive experimentations. In particular, we plan to study the possibility of map based exploration of similar symmetric transformations across shapes in near future. Another important direction that can be considered is the possibility of incorporating this technique in Computer Aided Geometric Design for urban architecture. In urban architecture, symmetric repetition of same pattern is a common thing and during the design phase, if the basic structure is stored only once and the symmetric repetitions are saved as *Functional Maps*, it can possibly solve both the space complexity problem and the design efficiency problem.

REFERENCES

- AUBRY, M., ULRICH, S., and DANIEL, C. (2011) The wave kernel signature: A quantum mechanical approach to shape analysis. *Proc. Computer Vision Workshops (ICCV Workshops)*.
- BOKELOH, M., BERNER, A., WAND, M., SEIDEL, H.-P., and SCHILLING, A. (2009) Symmetry detection using line features. *Computer Graphics Forum (Eurographics)* 28(2), pp. 697-706.
- BRONSTEIN, A. M., BRONSTEIN, M. M., BUSTOS, B., CASTELLANI, U., CRISANI, M., FALCIDIENO, B., GUIBAS, L. J., KOKKINOS, I., MURINO, V., SIPIRAN, I., OVSJANIKOV, M., PATANE, G., SPAGNUOLO, M., and SUN, J. (2010) SHREC 2010: Robust feature detection and description benchmark. *Proc. Workshop on 3D Object Retrieval (3DOR10)*.
- BRONSTEIN, A. M., BRONSTEIN, M. M., and KIMMEL, R. (2007) Calculus of non-rigid surfaces for geometry and texture manipulation. *IEEE Trans. Visualization and Computer Graphics*, Vol 13(5), September-October, pp. 902-913.
- BELKIN, M., SUN, J., and YUSU, W. (2008) Discrete Laplace operator on meshed surface. *Proc. Symp. Comp. Geometry*, pp. 278-287.



Fig. 12. Clustering of the functional maps resulting in three symmetry groups.

- BERNER, A., BOKELOH, M., WAND, M., SCHILLING, A., and SEIDEL, H.-P. (2008) A graph-based approach to symmetry detection. *Proc. Fifth Eurographics/IEEE VGTC Conf. Point-Based Graphics*, Eurographics Association, pp. 1–8.
- ELDAR, Y., LINDENBAUM, M., PORAT, M., and ZEEVI, Y. Y. (1997). The farthest point strategy for progressive image sampling. *IEEE Transactions on Image Processing*, 6(9), 1305-1315.
- IMIYA, A., UENO, T., and FERMIN, I. (1999). Symmetry detection by random sampling and voting process. *Proc. IEEE Conf. Image Analysis and Processing (CAIP99)*, pp. 400-405.
- KAZHDAN, M., CHAZELLE, B., DOBKIN, D., FUNKHOUSER, T., and RUSINKIEWICZ, S. (2003) A reflective symmetry descriptor for 3D models. *Algorithmica*, 38(1), October, pp. 201–225.
- LI, W., ZHANG, A., and KLEEMAN, L. (2005) Fast global reflectional symmetry detection for robotic grasping and visual tracking. *Proc. Australasian Conference on Robotics and Automation (ACRA05)*, December.
- LI, M., LANGBEIN, F., and MARTIN, R. (2006) Constructing regularity feature trees for solid models. *Geometric Modeling and Processing-GMP 2006*, Springer Berlin Heidelberg, pp. 267–286.
- LING, H., and OKADA, K. (2006) Diffusion distance for histogram comparison. *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR 2006)*, Vol. 1, pp. 246–253.
- LIU, S., MARTIN, R., LANGBEIN, F., and ROSIN, P. (2007) Segmenting periodic reliefs on triangle meshes. *Mathematics of Surfaces XII*, Springer Berlin Heidelberg, pp. 290-306.
- LIPMAN, Y., CHEN, X., DAUBECHIES, I., and FUNKHOUSER, T. (2010). Symmetry factored embedding and distance. *ACM Transactions on Graphics (TOG)*, 29(4), pp. 103.
- LIPMAN, Y., RUSTAMOV, R. M., AND FUNKHOUSER, T. A. (2010). Biharmonic distance. *ACM Transactions on Graphics (TOG)*, 29(3), pp. 27.
- LITMAN, R., and BRONSTEIN, A. (2013) Learning spectral descriptors for deformable shape correspondence. *IEEE Trans. Pattern Analysis and Machine Intelligence*, in press.
- MITRA, N. J., GUIBAS, L. J., and PAULY, M. (2006) Partial and approximate symmetry detection for 3D geometry. *ACM Transactions on Graphics (TOG)*, 25(3), pp. 560-568.
- MITRA, N. J., PAULY, M., WAND, M., and CEYLAN, D. (2013) Symmetry in 3D geometry: Extraction and applications. *Computer Graphics Forum*, Blackwell Publishing Ltd.
- OVSIJANIKOV, M., SUN, J., and GUIBAS, L. (2008) Global intrinsic symmetries of shapes. *Computer Graphics Forum* 27(5), Blackwell Publishing Ltd., pp. 1341-1348.
- OVSIJANIKOV, M., BEN-CHEN, M., SOLOMON, J., BUTSCHER, A., and GUIBAS, L. (2012). Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4), pp. 30.
- PAULY, M., MITRA, N. J., WALLNER, J., POTTMANN, H., and GUIBAS, L. (2008) Discovering structural regularity in 3D geometry. *ACM Transactions on Graphics (TOG)*, 27(3), pp. 43:1-11.
- PINKALL, U., and POLTHIER, K. (1993) Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2(1), pp. 15–36.
- PODOLAK, J., SHILANE, P., GOLOVINSKIY, A., RUSINKIEWICZ, S., and FUNKHOUSER, T. (2006) A planar-reflective symmetry transform for 3D shapes. *ACM Transactions on Graphics (TOG)* 25(3), July, pp. 549–559.
- PORKASS, J., BRONSTEIN, A. M., BRONSTEIN, M. M., SPRECHMANN, P., and SAPIRO, G. (2013) Sparse modeling of intrinsic correspondences. *Computer Graphics Forum*, 32(2.4), Blackwell Publishing Ltd., pp. 459–468.
- RAVIV, D., BRONSTEIN, A. M., BRONSTEIN, M. M., and KIMMEL, R. (2010) Full and partial symmetries of non-rigid shapes. *Intl. Jour. Computer Vision* 89(1), pp. 18-39.
- RAVIV, D., BRONSTEIN, A., BRONSTEIN, M., and KIMMEL, R. (2007) Symmetries of non-rigid shapes. *Proc. Intl. Conf. Computer Vision*, pp. 1–7.
- RUSTAMOV, R. (2008) Augmented planar reflective symmetry transform. *The Visual Computer* 24(6), pp. 423-433.
- SHIKHARE, D., BHAKAR, S., and MUDUR, S. (2001) Compression of large 3D engineering models using automatic discovery of repeating geometric features. *Signal Processing*, 19(20), pp. 233-240.
- WANG, Y., XU, K., LI, J., ZHANG, H., SHAMIR, A., LIU, L., CHENG, Z., and XIONG, Y. (2011) Symmetry hierarchy of man-made objects. *Computer Graphics Forum*, 30(2), Blackwell Publishing Ltd., pp. 287-296
- XU, K., ZHANG, H., TAGLIASACCHI, A., LIU, L., LI, G., MENG, M., and XIONG, Y. (2009) Partial intrinsic reflectional symmetry of 3D shapes. *ACM Transactions on Graphics (TOG)*, 28(5), pp. 138:1-10.
- XU, K., ZHANG, H., JIANG, W., DYER, R., CHENG, Z., LIU, L., and CHEN, B. (2012) Multi-scale partial intrinsic symmetry detection. *ACM Transactions on Graphics (TOG)*, 31(6), pp. 181.
- YEN, L., FOUSS, F., DECAESTECKER, C., FRANCO, P., and SAERENS, M. (2007) Graph nodes clustering based on the commute-time kernel. *Proc. 11th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD 2007)*, Springer Lecture notes in Computer Science, LNCS.
- YIP, R. (2000) A Hough transform technique for the detection of reflectional symmetry and skew-symmetry. *Pattern Recognition Letters*, 21(2), p. 117–130.